

FY09 Tactical Plan for **FermiGrid**

Prepared by: Keith Chadwick

Date: 19 September 2008

Relevant Strategic Plans - Strategic Plan for Grids 2009-2011 (SP4G)

SP4G.Goals and Objectives -

- Perform all work activities with awareness and understanding of appropriate safety practices and procedures.
- Operate, support, and evolve a robust, effective, secure, local production grid facility that supports the scientific program of Fermilab (FermiGrid).
- Increase FermiGrid's operational efficiency by the application of ITIL techniques, enhanced monitoring and alarming, and the creation of additional operational tools.
- Provide a grid services platform including services that are contributions to the OSG and help Fermilab achieve its goals. Special attention should be paid to implementations compatible with the infrastructure as outlined in the CMS Computing Technical Design Report (CTDR). Collaborate with external projects in the development and deployment of these grid services and grid tools.
- Understand and develop grid security policies, processes, and procedures in order to ensure the secure operation of the production grid facility. Assist the experiments/projects in developing operational security expertise.
- Work towards seamless interoperation between the national and regional grids that Run II and CMS worldwide collaborations depend on— in particular the EGEE (Enabling Grids for E-scienceE Project) and OSG.
- Work towards interoperation with additional national and regional grids such as TeraGrid and campus grids.
- Maintain visibility as a leading member of the OSG, used by all Fermilab scientific groups and take leadership roles within this consortium.
- Maintain and develop relationships and collaboration with sites, software contributors, committees and other bodies sufficient to increase the usability and administration of grid infrastructure integral to the Fermilab mission. Special attention should be paid to the needs of the DOE environment.
- Understand/Improve the relationship between VOs, resource providers, and Grid providers and the user experience working with grid tools.
- Offer excellent support to the Fermilab experiments/projects to enhance and enable their efficient use of FermiGrid, OSG, and world-wide grid facilities. Provide a support level that is appropriate to enable the virtual organization to operate in an efficient manner.

- Promote use of distributed computing facilities by additional experiments/projects including the neutrino and astrophysics programs as needed.
- Provide outreach and communication activities to further the movement of Fermilab experiments/projects to using the distributed grid infrastructure.
- Investigate the application of MPI technologies in the grid environment. Deploy these technologies in the grid environment to support the MPI applications.
- Investigate the utility of cloud computing to the grid community and apply to distributed computing at Fermilab as appropriate.
- Participate in the identification of a grid/distributed computing based path forward for HEP and astrophysics communities.

SP4G.Strategies –

- Work with specific experiments/projects to both gain domain knowledge and foster acceptance of the distributed computing approach.
- Use CMS and the Run II experiments as exemplars and critical leverage to evolve and expand distributed computing techniques and technologies.
- Undertake leadership and supporting roles in various organizations, including the Open Science Grid, and evolving the organization to build acumen in applicable techniques.
- Partner and collaborate with other significant contributors in developing and deploying distributed computing middleware and production facility capabilities.
- Provide common grid services at Fermilab, including the operation of the FermiGrid campus grid.
- Evolve site operational strengths to a distributed computing context.
- Work with other departments in CD to facilitate the deployment of grid software and its operation and support.

FermiGrid.Accomplishments during FY08 –

The FermiGrid accomplishments during FY08 include:

1. Throughout FY2008, the FermiGrid infrastructure has been maintained and extended, adding support for new services and increasing the level of inspection of the Grid middleware services and the underlying network and system components. The highest observed rate of Grid User Mapping Service (GUMS) requests for the period 01-Jul-2007 through 30-Jun-2008 was 1,028,388 service requests/day on 11-Oct-2007.
2. Based on the prior observed growth in the service requests rate, a project to design and implement a highly available “FermiGrid-HA” system was undertaken as part of the FY07 through FY08 FermiGrid effort. The project scope was to support highly available operations (allow for continued operation in the event of hardware, network or software failures together with additional capacity) for the FermiGrid Virtual Organization Management Service (VOMS), Grid User Mapping Service (GUMS) and Site AuthoriZation (SAZ) services.
3. The design of FermiGrid-HA utilized three technologies: Hardware virtualization (using the Xen hypervisor with Scientific Linux), MySQL circular database replication, and Linux Virtual Server (LVS) load balancing (to automatically distribute service requests across the pool of services and transparently remove failed systems from the service pool).
4. The hardware procurement necessary to support the FermiGrid-HA design was completed in July 2007 and the software installation and configuration was performed over the July 2007 through October 2007 period.
5. During the month of November 2007, as part of the final FermiGrid-HA commissioning work a series of load and stress tests were performed against the new infrastructure. The results of the load and stress tests showed that the new “FermiGrid-HA” system was capable of supporting in excess of 10,000,000 service requests/day (>115 requests/second).
6. FermiGrid-HA formally commenced production operations for VOMS-HA, GUMS-HA and SAZ-HA on 01-Dec-2007. For the period of 01-Dec-2007 through 30-Jun-2008, the total downtime for FermiGrid-HA services was 10 minutes, resulting in an overall service availability of 99.9967%.
7. An introductory set of courses was developed for the “FermiGrid School”: FermiGrid 101 – FermiGrid Introduction and Overview, FermiGrid 201 – Scripting and Running Grid Jobs and FermiGrid 202 – Grid Storage Access. These courses were well received and offered a total of four times during spring to early summer 2008.

8. Following the production deployment of FermiGrid-HA, a review of the other internal and externally visible services (Squid, Ganglia, Zabbix, Syslog-Ng, Information Gatherer, etc.) offered by FermiGrid was performed in order to determine the applicability of similar “HA” techniques to these services. Based on this review, a plan was developed to transition additional services to “HA” services. At the present time, the additional service transition plan is well underway, with the majority of the services migrated to “HA” or “HA-like” configurations. Work is currently underway on the necessary underlying technologies that will be required to complete the Gatekeeper-HA and MyProxy-HA deployments.
9. FermiGrid personnel made significant contributions to the Fermilab Kerberos Certificate Authority (KCA) Distinguished Name (DN) transition. Starting with performing the necessary research and tests to verify proper processing of the new proposed formats for KCA DNs across FermiGrid and the OSG. Once the details of the transition was finalized, FermiGrid arranged for the population of the new format KCA DNs into all of the VOMRS and VOMS servers managed by FermiGrid. And two weeks following the transition, removing the old format KCA DNs from all of the VOMRS and VOMS servers managed by FermiGrid. This sequence of operations resulted in a transparent migration from the old format KCA DNs to the new format KCA DNs for the Grid users.
10. FermiGrid personnel made significant contributions to the Fermilab Open Science Enclave (OSE) both within the formal Open Science Enclave Working Group (OSE-WG) and during the day-to-day FermiGrid support and operations. These contributions include updates to the OSE security plan, risk assessment and contingency plans, development of the OSE Baseline (currently in draft form), developing mechanisms to enumerate the list of systems that are in the OSE, and other contributions.
11. FermiGrid also performed significant upgrades to the Fermilab Site AuthoriZation (SAZ) service. These upgrades included development of an exhaustive regression test suite to validate new releases of SAZ, making necessary bug fixes in the SAZ source, modification of the SAZ server code to implementing hibernate database access, development of a formal SAZ project plan, ...
12. In June 2008, the responsibility for production OSG and Fermilab Gratia accounting repository operations was formally transferred to FermiGrid. Prior to this transition FermiGrid had been responsible for only system support issues.
13. During early CY2008, FermiGrid personnel developed the technology to deploy a transparent OSG to TeraGrid gateway. This gateway was formally opened to a small test of test users during July 2008.
14. During the month of September 2008, FermiGrid personnel performed a complete hardware “refresh” of the production systems, and started the hardware refresh of the FermiGrid test-bed systems (using the former production systems).

15. In addition to the “core” FermiGrid services work above, members of FermiGrid also maintained the CDF, D0 and GP Grid cluster Globus Gatekeeper systems (fedfosg[1-4], d0cabosg[1-2], fnpc[3,4]x1) and corresponding batch management systems (fedfcm[1-3], fnpc5x1).

FermiGrid.Tactical Objectives for FY09 –

The tactical objectives for FermiGrid effort in FY09 are:

1. Operate the FermiGrid services at a very high level of reliability, with minimal disruptions to normal operations together with quick resolution of problems and trouble tickets. Integrate ITIL processes into FermiGrid operations as appropriate. Continue to enhance FermiGrid metrics and service monitoring capabilities and infrastructure.
2. Maintain and extend the FermiGrid-HA deployment. The scope of this objective includes routine expansion and upgrades of the existing infrastructure (hardware and software) to keep pace with the support needs from both internal FermiGrid clients (CMS, CDF, D0, GP grid, etc.) and clients from the Open Science Grid, WLCG, etc.
3. Maintain and extend the associated OSG Grid services infrastructure (OSG and Fermilab Gratia accounting systems, OSG and Dzero the Resource Selection Service, OSG VOMRS/VOMS). The scope of this objective includes routine expansion and upgrades of the existing infrastructure (hardware and software) to keep pace with the support needs from both internal FermiGrid clients and clients from the Open Science Grid, WLCG and possibly TeraGrid.
 - a. The production Fermilab and OSG Gratia accounting repository systems are nearing the end of their warranty period. In addition, these systems need significant revisions to their hardware deployment to support both HA (high availability) and HP (high performance) operations. A total of \$60K (operating) is necessary to accomplish this work. This can be funded in two stages (of \$30K each) with stage 1 being the upgrade of the Fermilab Gratia accounting repository and stage 2 being the upgrade of the OSG Gratia accounting repository.
 - b. The production and integration ReSS systems are close to the end of their warranty period, and the hardware is showing signs of impending failures. In addition, the ReSS service and systems need revisions to their deployment configuration in order to support a HA (high availability) deployment. These configuration revisions are well understood. A total of \$20K (operating) is necessary to accomplish this work.
4. Provide necessary support to the Laboratory’s computer security efforts, including incident handling effort and enhancements to security related tools. The FermiGrid operations personnel have recently been commissioned to serve as “expert” advisors to FCIRT on Grid issues.

5. Continue the development of the FermiGrid Site AuthoriZation (SAZ) service. The scope of this objective includes routine maintenance and upgrades of the SAZ service as well as working with Open Science Grid and Globus to extend SAZ capabilities to meet the needs of these communities. A significant milestone will be receipt of the formal technical requirements for the “banning tool” from both the Open Science Grid and Globus communities.
6. Continue the ongoing support of the Fermilab Grid cluster gatekeepers that are administered by FermiGrid personnel. Deploy “Gatekeeper-HA” functionality on these gatekeepers.
7. Maintain and enhance the support for the existing FermiGrid user and stakeholder communities, provide hosts for submission of grid jobs, provide development hosts for development of grid software. Undertake outreach initiatives to bring on additional FermiGrid user and stakeholder communities onboard. Including the further development of the FermiGrid “Grid School” as well as ongoing travel and collaboration with current and potential stakeholder communities. A critical component for the sustainability of this objective is the FermiGrid Cloud Computing Initiative (item 9 below).
8. Deploy an initial Grid MPI capability for selected Fermilab Grid users and for opportunistic Open Science Grid users. Investigate with the user community sustainable methods for MPI application development and deployment in a Grid environment. A total of \$35K (operating) is requested to acquire the required Gigabit Ethernet switch, Infiniband switch and interfaces.
9. Deployment of a FermiGrid Cloud Computing infrastructure to replace the obsolete systems in the GridWorks/FAPL cluster as well as provide a sustainable mechanism for provisioning of custom development environments for CD personnel. Collaborate with CD Departments and external researchers to investigate and deploy sustainable mechanisms for virtual machine configuration, deployment, and scheduling. A total of \$240K (equipment) and \$10K (operating) is requested to replace the existing GridWorks/FAPL systems.
10. Participation in the Open Science Grid (OSG) development, integration, deployment and operations activities.

FermiGrid.Activity Tasks for FY09 –

The proposed activity tasks for FermiGrid effort in FY09 are:

Grid / FermiGrid / Operations

- Description: FermiGrid Services Day-To-Day Operations including High Availability
- Activity type: Service (ongoing)
- Timescale: Continuous
- Objectives: Maintain FermiGrid infrastructure.
Respond to user problems.
- Milestones: n/a – Ongoing operation.
- Metrics: Successful use of the FermiGrid infrastructure by Fermilab personnel, OSG collaborators and LHC experiments. Number of (automatic and manual) interventions required to maintain the infrastructure availability at Service Levels specified in formal SLA's. Collection and publication of individual service metrics (see <http://fermigrid.fnal.gov/fermigrid-metrics.html>).

Grid / FermiGrid / Development

- Description: FermiGrid Services Development
- Activity type: Project (ongoing)
- Timescale: Continuous
- Objectives: Ongoing development of the FermiGrid infrastructure and Grid middleware.
- Milestones: Receipt of the formal technical requirements for the “banning tool” from both the Open Science Grid and Globus communities. New releases of software and Grid middleware.
- Metrics: n/a.

Grid / FermiGrid / Integration

- Description: FermiGrid Services Integration
- Activity type: Service (ongoing)
- Timescale: Continuous
- Objectives: Timely integration of newly released Grid middleware into the FermiGrid infrastructure.
- Milestones: Completion of the Grid middleware and other software integration cycles (following the FermiGrid Software Acceptance Process as documented in CD DocDB 2684-v4).
- Metrics: Number of interventions required to deploy newly released software into production operations.

Grid / FermiGrid / Gratia

- Description: FermiGrid Services Gratia Operations
- Activity type: Service (ongoing)
- Timescale: Continuous
- Objectives: Ongoing operation of the OSG and Fermilab Gratia infrastructure.
- Milestones: Production deployment of new hardware to host the Fermilab and OSG Gratia repositories. Timely deployment of new Gratia software releases (following the FermiGrid Software Acceptance Process as documented in CD DocDB 2684-v4).
- Metrics: n/a.

Grid / FermiGrid / Management

- Description: FermiGrid Services Management
- Activity type: Project (ongoing)
- Timescale: Continuous
- Objectives: Ongoing management of the FermiGrid Services project.
- Milestones: Timely deployment of new software releases (following the FermiGrid Software Acceptance Process as documented in CD DocDB 2684-v4).
- Metrics: n/a.

Grid / FermiGrid / User Support

- Description: FermiGrid Services User Support and Outreach.
- Activity type: Service (ongoing)
- Timescale: Continuous
- Objectives: Ongoing support for the existing FermiGrid user community and new community outreach.
- Milestones: n/a.
- Metrics: User satisfaction, number of issues resolved.

Grid / FermiGrid / Cloud Computing

- Description: FermiGrid Cloud Computing
- Activity type: Project (new initiative and development)
- Timescale: Continuous
- Objectives: Production deployment of a Cloud Computing Infrastructure within FermiGrid.
- Milestones: Commissioning of the FermiGrid Cloud.
- Metrics: User satisfaction, number of (configured, running) virtual machines.

Grid / FermiGrid / MPI

- Description: FermiGrid MPI
- Activity type: Project (new initiative and development)
- Timescale: Continuous
- Objectives: Production deployment of an MPI infrastructure within FermiGrid.
- Milestones: Commissioning of the GPMPI cluster
- Metrics: User satisfaction, number of MPI jobs.

FermiGrid.Priorities for FY09:

1. Continued operation of the FermiGrid Services infrastructure. Continued maintenance and extension of FermiGrid-HA. Continue to enhance FermiGrid metrics and service monitoring capabilities and infrastructure. Extend the operational model of the FermiGrid infrastructure to a 24x7 operation if possible.
2. Extend FermiGrid-HA deployment.
3. Ongoing support for the services operated for the Open Science Grid (OSG VOMS/VOMRS, OSG Gratia, OSG to TeraGrid gateway, etc.).
4. Provide the necessary support to the Laboratory's computer security efforts, including incident handling effort and enhancements to security related tools.
5. Continue the development of the FermiGrid Site AuthoriZation (SAZ) service.
6. Ongoing support of the Fermilab Grid cluster gatekeepers that are administered by FermiGrid personnel.
7. Respond to and support the existing FermiGrid user and stakeholder communities, provide development hosts for development of grid software. Undertake outreach initiatives to bring on additional FermiGrid user and stakeholder communities onboard. Continue the development of FermiGrid "Grid School" courses as well as ongoing travel and collaboration with current and potential stakeholder communities.
8. Deployment of a limited general purpose MPI capability. The initial investigation shall take place using ~24 nodes which were purchased as part of the FY07 GP Grid cluster expansion.
9. Deployment of an initial FermiGrid Cloud Computing infrastructure.
10. Continued participation (including ongoing travel) in the Open Science Grid (OSG) development, integration, deployment and operations activities.

If we are short of time or resources, then we will postpone, eliminate, or reduce participation in the lower priority activities.

FermiGrid.Staffing for FY09:

Progress on FermiGrid is significantly effort limited. The FermiGrid staffing was reduced by 1 FTE in January 2008 with the retirement of Valery Sergeev. Due to the FY08 budget crisis, this effort was not replaced. The 1.8 new FTEs that were authorized as part of the initial FY08 budget were also eliminated as part of the response to the FY08 budget crisis.

Significant additional operational (Gratia) and development (SAZ) activities were added to FermiGrid during FY08. For FY09, in addition to the new activities added in FY08, significant additional development, integration and operational activities (SAZ, GPMPI, Cloud Computing) are included in the FermiGrid FY09 Tactical Plan. Additional personnel to contribute to these activities would greatly enhance the ability of the FermiGrid to meet the internal milestones as well as the external (OSG) milestones. FermiGrid does not currently have sufficient staff to successfully implement a 24x7 on-call pager rotation or to extend significant operational support for other services. Without at least one (1) and preferably two (2) additional staff, FermiGrid will be unable to offer significant operational support for additional services or attempt a 24x7 on-call pager rotation.

Staff	Assignment	Effort
Berman, Eileen	Grid Dept Line Management	40%
Chadwick, Keith	Grid Dept Line Management	20%
	FermiGrid Project Manager	20%
	FermiGrid Operations, Development, etc.	50%
	Security Activities	10%
Sharma, Neha	FermiGrid Development	50%
Timm, Steve	Assistant Group Leader	5%
	FermiGrid Operations, Development, etc.	95%
Yocum, Daniel	FermiGrid Operations, Development, etc.	25-50%
	Gratia Operations	25-50%
	GridWorks cluster coordinator / Cloud Computing	Up to 25%
Green, Chris	Gratia Operations	25%
*** New Hire 1	FermiGrid Operations, Development, etc.	100%
*** New Hire 2	FermiGrid Operations, Development, etc.	100%

FermiGrid.Change Control for FY09:

If line management requires a change to the FermiGrid deliverables, or we find that we have inadequate resources within our budget we will file a formal change request to this plan. Any change request must originate with the FermiGrid Project Leader (Keith Chadwick) and proceed through the corresponding Department and Division Channels. The stakeholders who need to be informed of any change are the corresponding FermiGrid Stakeholder(s) (OSG, CDF, D0, CMS, GP Grid, etc.) together with both Department and Division management.

FermiGrid.Risk Assessment for FY09:

1. Failure to insure continued operation of the FermiGrid infrastructure may result in significant production service outages for both the multiple clusters on the Fermilab site that depend on the services (Site Globus Gateway, VOMS, GUMS, SAZ and MyProxy) provided by the core FermiGrid infrastructure. Failure to extend the Metrics and Monitoring infrastructure may result in service interruptions and/or necessary capacity planning information to fail to be available. Failure to extend the operational model of the FermiGrid infrastructure to a 24x7 operation may also result in significant production service outages for both the multiple clusters on the Fermilab site that depend on the services.
2. Failure to maintain and extend the FermiGrid-HA deployment may result in significant production service outages.
3. Failure to maintain and extend the support for the services operated for the Open Science Grid may result in poor relations with the corresponding OSG stakeholders and may result in the services being relocated to other sites. Failure to acquire the proposed hardware may result in outages of the corresponding service, or the service not being able to handle the increased loads that are expected during FY09 and following years.
4. Failure to provide the necessary support to the Laboratory's computer security efforts (incident handling effort and enhancements to security related tools) may result in potential security vulnerabilities not being addressed adequately or the response capabilities of FermiGrid being insufficient to respond to an actual Grid security incident.
5. Failure to maintain and extend the FermiGrid Site AuthoriZation (SAZ) service may result in FermiGrid being at risk during an actual Grid security incident. Failure to work with the Open Science Grid and Globus communities to insure the adoption of SAZ as the official Grid wide banning tool may result in these communities developing or selecting incompatible banning tools which would result in an increased support load on FermiGrid personnel to maintain both SAZ and the external banning tool (if the external banning tool did not offer feature for feature compatibility with SAZ).

6. Failure to provide the necessary support to the cluster gatekeepers may result in increased downtime or service interruptions and corresponding outages of the experiments offline analysis capabilities.
7. Failure to provide the necessary support to existing Grid users and stakeholders and outreach initiatives to bring on additional FermiGrid user and stakeholder communities onboard will result in significant “ill-will” between the CD and the growing grid communities and may result in both FermiGrid and CD becoming “niche” players in the global grid effort. Failure to undertake the necessary ongoing travel and collaboration will significantly interfere with FermiGrid participation in the global grid effort.
8. Failure to deploy a general purpose MPI capability will result in the existing and new groups or organizations that need MPI capabilities not being well served by existing Grid installations at Fermilab.
9. Failure to deploy a Cloud computing infrastructure to replace the obsolete systems in the GridWorks/FAPL cluster will result in power and cooling issues as well as continue the ongoing lack of resources for development.
10. Failure to effectively participate in the Open Science Grid (OSG) development, integration, deployment and operations will result in FermiGrid falling behind the curve of grid software development and deployment. If that happens, the grid community will bypass Fermilab, or significant additional effort will be needed on an emergency basis to bring FermiGrid back up to the level expected by the grid community.